MEMORANDUM

No 09/2003

**How Tight is the Labour Market?
A Micro-Based Macro Indicator**

By
Simen Gaure and Knut Røed

ISSN: 0801-1117

Department of Economics
University of Oslo

This series is published by the **University of Oslo** **Department of Economics** P. O.Box 1095 Blindern N-0317 OSLO Norway Telephone: + 47 22855127 Fax: + 47 22855035 Internet: http://www.oekonomi.uio.no/ e-mail: econdep@econ.uio.no	In co-operation with **The Frisch Centre for Economic Research** Gaustadalleén 21 N-0371 OSLO Norway Telephone: +47 22 95 88 20 Fax: +47 22 95 88 25 Internet: http://www.frisch.uio.no/ e-mail: frisch@frisch.uio.no

List of the last 10 Memoranda:

No 08	Finn R. Førsund and Nikias Sarafoglou The Tale of two Research Communities: The Diffusion of Research on Productive Efficiency. 43 pp.
No 07	Taryn Ann Galloway and Rolf Aaberge Assimilation Effects on Poverty Among Immigrants in Norway. 26 pp.
No 06	Geir B. Asheim Green national accounting with a changing population. 23 pp.
No 05	Geir B. Asheim and Ylva Søvik The semantics of preference-based belief operators. 30 pp.Rolf
No 04	Zhiyang Jia A Mixture Model of Household Retirement Choice. 21 pp.
No 03	Erling Eide Optimal Provision of Public Goods with Rank Dependent Expected. tility. 21 pp.
No 02	Hilde C. Bjørnland Estimating the equilibrium real exchange rate in Venezuela. pp.
No 01	Svenn-Erik Mamelund Can the Spanish Influenza pandemic of 1918 explain the baby-boom of 1920 in neutral Norway?. 33 pp.
No 36	Elin Halvorsen A Cohort Analysis of Household Saving in Norway. 39 pp.
No 35	V. Bhaskar and Steinar Holden Wage Differentiation via Subsidised General Training. 24 pp.

A complete list of this memo-series is available in a PDF® format at:
http://www.oekonomi.uio.no/memo/

4 March 2003

How Tight is the Labour Market? A Micro-Based Macro Indicator*

Simen Gaure
Centre for Information Technology Services, University of Oslo and
The Ragnar Frisch Centre for Economic Research, Oslo

and

Knut Røed
The Ragnar Frisch Centre for Economic Research, Oslo

Abstract

We develop a new indicator of labour market tightness, based on the pure calendar time changes in individuals' transition rates from unemployment to employment. Based on Norwegian register data from the 1989-2002 period, we show that this indicator, in contrast to the aggregate rate of unemployment, correlates well with an ex-post-calculated GDP-based business cycle indicator, even around the time of business cycle turning points. The indicator can be calculated just as quickly as the unemployment rate, both at an aggregate and a disaggregate level, and hence improve policy makers ability to assess current labour market developments.

Keywords: Labour market tightness, Business cycles, Unemployment

JEL Classification: C41, E32, J64.

* We wish to thank Rolf Aaberge, Morten Petter Johansen, Per Richard Johansen and Steinar Strøm for valuable comments, and the Norwegian Directorate of Labour and the Norwegian Ministry of Finance for financial support. Correspondence to: Knut Røed, Gaustadalléen 21, 0349 Oslo. E-mail: knut.roed@frisch.uio.no.

1 Introduction

Updated knowledge about the state of the labour market is of vital importance for monetary and fiscal policy makers. Yet there exist no easily observed aggregates that both timely and reliably gauge its tightness. Policy makers typically keep a close eye on frequently updated aggregates, such as the rate of unemployment. In this paper, we argue that changes in the rate of unemployment in some cases give a misleading picture of labour market developments as seen from a given agent in a given labour market position. In particular it tends to display its troughs and peaks long after the business cycle has actually turned. There are two reasons for this. The first reason is simply that it takes time before changes in individuals' transition probabilities level out the flows into and out of the unemployment pool. The unemployment rate will increase (decrease) as long as inflow exceeds (falls short of) outflow, irrespective of improvements (deteriorations) in the underlying transition probabilities. The second is that the time it takes to level out the flows depends on the *composition* of the stocks in terms of individual employability. In particular, it depends on the composition of the unemployment pool, which tends to vary systematically over the business cycle.

The selection-issue is of course well known to micro-econometricians, and it plays a key role in econometric analyses of unemployment durations. However, there has been little scope for transferring this knowledge from the micro to the macro level in a way that can shed light on business cycle developments. With the emergence of large-scale register data, this may change. To an increasing extent, micro-economists will gain access to micro data that adds up to the aggregates used by the macro-economists (Røed and Raaum, 2003). This facilitates a proper decomposition of aggregate developments into components related to external developments faced by each unit and developments related to changes in the composition of units within the aggregate.

In the present paper, we conduct this kind of decomposition with respect to the outflow rate from the unemployment pool. The idea is to identify the pure *calendar time* component within the context of a hazard rate model, along the lines suggested by Imbens and Lynch (1993). The *business cycle* is then interpreted as the trend-wise (or smoothed) change over time in individuals' probability of escaping unemployment for a job, ceteris paribus. Hence, there are two distinct steps involved in the computation. The first step is to identify and estimate the 'effect' of calendar time on job transitions, conditioned on everything else (observed characteristics, unobserved characteristics, spell duration). We do this in a non-parametric fashion. The second step is to decompose the estimated calendar time effects into business cycles, seasonal cycles and irregular components. At this point, we rely on standard decomposition techniques, i.e. X12-ARIMA (with appropriate 'trading day' adjustments).

One possible objection to the business cycle interpretation of 'pure' calendar time changes in employment hazards is that changes in employment prospects may also induce changes in the reservation wages (or choosiness) of the job seekers. Hence, improved job prospects do not necessarily result in higher hazard rates. This is a valid theoretical argument, although it has been proved that the job offer arrival rate and the hazard rate do move in the same direction under reasonable assumptions[1] regarding the wage distribution (Burdett and Ondrich, 1985). Adjustments in the reservation wage cushion the impact of changes in labour demand on the transition rate into jobs, but do not alter the fundamental comonotonicity. It may also be noted that a

[1] A sufficient condition is that the left-truncated mean of the wage offer distribution has a slope less than one with respect to the truncation point (the reservation wage). A sufficient condition for this to be satisfied is that the densities of the wage offer distributions are log-concave (see e.g. An, 1998).

functional relationship between job choosiness and the underlying business cycle will affect the interpretation of virtually any conceivable business cycle indicator (any indicator that is affected by employment decisions). This implies in particular that the common practice of assessing the relationship between the hazard rate and the business cycle by means of using a *rate of unemployment* as the explanatory business cycle indicator is disputable. If calendar time variations in the hazard rates out of unemployment do not reflect changes in labour demand, there is no reason to believe that variations in its stock can be expected to do so. The stock of unemployment is of course affected by everything that affects the flows into and out of unemployment.

The concepts of 'business cycle' and 'labour market tightness' do not conform to generally accepted model-based definitions. While the existing empirical literature has typically discussed the cyclical behaviour of the employment hazard rate on the (implicit) presumption that the true business cycle is captured by either the rate of unemployment or the rate of GDP growth (see e.g. Bover et al, 2002, and references therein), we instead characterize the business cycle (labour market tightness) in terms of the 'pure calendar time changes' in individual employment hazards. And on the basis of this characterization, we show that the aggregate rate of unemployment tends to behave *procyclically* around the time of business cycle turning points. Modern theories of wage formation (such as bargaining models and efficiency wage models) suggest that the wage pressure is determined by the development in individual employment transition probabilities (the outside options), and not by the stock of unemployment (see e.g. Layard et al., 1991, p. 145). Hence, from the viewpoint of fiscal and monetary policy makers, it is important to be able to observe the development of these transition probabilities as quickly and precisely as possible.

For the purpose of illustration, we have gathered register data containing all unemployment spells in Norway starting between March 1989 and July 2002, and we use these data to estimate a monthly business cycle indicator. We put the selection dynamics on display by decomposing the correlation between predicted employment transition rates on the one hand, and calendar time and spell duration on the other, into the two driving mechanisms of *selection* and *causality*. It turns out that selection forces play a relatively modest role in explaining pure calendar time changes in average transition rates, while they play a key role in explaining the spell duration changes.

Although the information content in the data is sufficient for non-parametric identification of the roles played by calendar time, process time (spell duration) and observed and unobserved heterogeneity, actual estimation raises huge computational problems. In order to ripe the full benefit from the richness of register data, it is essential to avoid unjustified parametric restrictions that potentially produce unpredictable biases in the parameters of interest. Some structure must be imposed on the data, however, otherwise the parameters always outnumber the observations, and nothing of interest can be extracted from them. Our paper serves two purposes. The first is to contribute to a better understanding of labour market business cycles and thereby enhance policy makers' ability to identify turning points as quickly as possible. The second is to present a general method for disentangling the different mechanisms that produce changes in hazard rates over time, such as duration dependence, unobserved heterogeneity and business and seasonal cycles – a method that is applicable even in situations with millions of observations and thousands of unknown parameters and hence able to utilize register data in an efficient way. In the next section, we outline our statistical tool, in the form of a discrete hazard rate model with calendar time effects. Section 3 presents an application of the model to Norwegian register data covering the 1989-2002 period. Section 4 concludes. An efficient way of estimating the type of models that we use in the paper is presented in an Appendix.

2 The Statistical Model

Let D_i be the stochastic duration of an unemployment spell for individual i. Let t indicate calendar time and let $\overline{t}_i$ be the point in time at which the spell started, such that $(t-\overline{t}_i)$ is the spell duration at calendar time t. Assume first that the employment hazard rate can be decomposed into one factor that depends on calendar time and another that depends on everything else. This is the *homogenous cycle assumption*, which implies that the hazard rate $\theta(i,t,t-\overline{t}_i)$ is multiplicatively separable can be expressed as follows

$$\theta(i,t,t-\overline{t}_i) \equiv \lim_{\Delta t \to 0} \frac{P\left(t-\overline{t}_i < D_i < t-\overline{t}_i + \Delta t \mid t-\overline{t}_i < D_i\right)}{\Delta t} = b(t)h(x_{it}, t-\overline{t}_i, v_i), \qquad (1)$$

where x_{it} is a vector of observed (potentially time-varying) variables and v_i is an unobserved individual (time-invariant) characteristic. The factor of interest is *b(t)* i.e. the pure calendar time effects, but in order to identify these properly it is essential to maintain *h(.)* as flexible as possible so as to prevent missing variables or invalid restrictions (that could conceivably have varying effects over time) from contaminating estimates of *b(t)*. While (1) is formulated in continuous time, data are typically discrete. We assume without loss of generality that the unemployment state is monitored at each integer point in time. Accordingly, the probability of observing an exit between two observation points (*t-1*) and t, conditional on being at risk the first of these observation points, is given by

$$P\left(t-\overline{t}_i - 1 < D_i < t-\overline{t}_i \mid t-\overline{t}_i - 1 < D_i\right) = 1 - \exp\left(-\int_{t-1}^{t} b(u)h(x_{it}, u-\overline{t}_i, v_i)du\right). \qquad (2)$$

For the sake of analytical tractability, we assume that time varying covariates do not change within each time unit. Moreover, we assume that spell duration effects are constant within each time unit. We then have that

$$\int_{t-1}^{t} b(u)h(x_{it},u-\overline{t}_i,v_i)du = h(x_{it},t-\overline{t}_i,v_i)\exp(\sigma_t), \quad \sigma_t \equiv \log\int_{t-1}^{t} b(u)du, \tag{3}$$

hence without loss of generality, we let the calendar time effect associated with each time interval occurring in the data be represented by its own period-specific parameter.

A further generalisation to an *idiosyncratic cycle model* is straightforward. Let *g=1,2,...,G* denote groups of observation types suspected to be affected differently by business or seasonal cycles. The idiosyncratic model then amounts to estimate separate calendar time parameters for each of these groups, i.e. σ_{gt}. In practice, this is done by introducing interaction terms between calendar time indicator variables and other explanatory variables in the model. The grouping may either be determined on the basis of prior knowledge or interest, or through a model reduction exercise. To the extent that the grouping is fixed at the individual level, idiosyncratic cycle models can be produced simply by dividing the population into separate datasets and then perform separate estimations. But if individuals switch between different groups during unemployment spells (or from one spell to another), a simultaneous estimation seems preferable. In particular, a simultaneous model is required in order to estimate separate business cycle indicators for different unemployment spell durations.

We specify the factor of proportionality *h(.)* as a flexible function of individual characteristics and spell duration, i.e. in the form of a separate dummy for each possible duration, a separate dummy for each possible age and so forth, in addition to a number of interaction terms, i.e. $h(x_{it},t-\overline{t}_i,v_i) = \exp\left(x_{it}^{'}\beta + \lambda_{t-\overline{t}_i} + x_{it}^{*'}(t-\overline{t}_i)\gamma + v_i\right)$, where the vector of explanatory variables x_{it} consists of a large number of dummy variables and $x_{it}^{*'}$ is a subset of these variables organised cardinally (i.e., while e.g. age appears as a set of dummy variables, one for each age measured in years, in x_{it}, it appears as a single scalar variable (age) in x_{it}^{*}). The parameters $\lambda_{t-\bar{t}_i}$ are the duration-month-specific effects (i.e. the parameters attached to the spell duration dummy variables). Unobserved heterogeneity is also modelled non-parametrically by assuming that the variables v_i are discretely distributed (Lindsay, 1983), with the number of mass-points chosen by adding points - one by one - until it is no longer possible to increase the likelihood function any further (Heckman and Singer, 1984). An important point to note is that the effects of unobserved heterogeneity are *identified from the data alone*, i.e. we do not have to rest on any of the model assumptions, apart from the assumption that the unobserved covariate is fixed at the individual level. There are two sources of identification. The first is the existence of 'lagged' variation in hazard rates for all durations above zero, conditional on the 'current' hazard rates. This kind of variation is ascertained through the inclusion of several cohorts of unemployed who face different business -and seasonal cycle conditions during their spells[2]. Identification of unobserved heterogeneity based on variation in lagged explanatory variables is more thoroughly discussed in Røed and Zhang (2003), and Brinch (2000) provides a formal proof for the proposition that it is sufficient for identification, even in the absence of a proportional hazard model. The second source of identification is the existence of repeat spells experienced by the same individuals (Van den Berg, 2001). To the extent that one is ready to assume that there are no causal linkages between such repeat spells, and that the individual unobserved covariates are constant across these spells, this introduces a kind of 'fixed effect' element in the identification strategy.

[2] The same source of identification is used by Abbring et al (2002) to disentangle the roles of heterogeneity and duration dependence from *aggregate* flows out of unemployment.

Let B_i be the number of spells experienced by individual i during the whole estimation period. Let t_{ib} be the calendar month in which spell number b ended and let $y_{ib}=1$ if the spell ended with a transition (non-censored), and zero otherwise. Finally, let W be the number of mass points in the distribution of unobservables. The likelihood function corresponding to the *homogenous cycle model*, in terms of observations of $(t_{ib}, \overline{t}_{ib}, y_{ib}, x_{it})$, is then given as

$$
\begin{gathered}
L_I = \prod_{i=1}^{N_I} \sum_{w=1}^{W} p_w \prod_{b=1}^{B_i} \left(\left(\varphi(t_{ib}, t_{ib} - \overline{t}_{ib}, x_{it}, v_w) \right)^{y_{ib}} \prod_{s=\overline{t}_i+1}^{t_{ib}-y_{ib}} \left(1 - \varphi(s, s - \overline{t}_{ib}, x_{is}, v_w)) \right) \right), \quad \sum p_w = 1, \\
\varphi(t, t-\overline{t}, x_{it}, v_i) = 1 - \exp(-\exp\left(x_{it}^{'}\beta + \lambda_{t-\overline{t}} + x^{*'}_{it} \log(t-\overline{t})\gamma + v_i + \sigma_t \right),
\end{gathered} \tag{4}
$$

where N_I is the number of individuals in the dataset, and p_w is the probability that an individual is characterised by and unobserved variable equal to v_w. This likelihood is maximised with respect to $\left(\sigma_t, \lambda_{t-\overline{t}}, \beta, \gamma, W, p_w, v_w\right)$.

Maximisation of the function (4) is well known to be difficult even with relatively few observations and few parameters. The reason is that the likelihood function is not globally concave, and that it may be quite flat in large areas around the local maximum points. In the context of our model, applications may involve millions of observations and thousands of parameters. And since the computational cost typically grows with the square of the number of parameters, the problem quickly becomes intractable in the sense that it exceeds accessible computational capacity by several orders of magnitude. The problem could of course be 'solved' by imposing restrictions on the model, i.e. reduce the number of free parameters. However, that would imply that the richness of the data was not fully exploited, and unpredictable biases could arise. In order to be sure that identification of calendar time and spell duration effects is based on data only, we consider the non-parametric approach to be of essential importance. We have therefore developed an optimisation routine that is tailored to non-parametric models, i.e. models in which all (or most) of the explanatory variables are dummy coded. This program builds on the concept of 'implicit dummy variables', which in essence reduce any number of mutually exclusive indicator variables to a single variable. In addition it builds on a very efficient optimisation procedure and facilitates the use of several computers (CPU's) simultaneously. The program is more thoroughly described in a separate Appendix.

3 An Illustrative Application: Norwegian Business Cycles in the 1990's

We applied the models described in Section 2 on Norwegian register data encompassing all unemployment spells in Norway starting between March 1989 and July 2002. In order to be sure that exits from the unemployment pool could safely be interpreted as entries into jobs, we focused on insured unemployment spells[3] for persons aged 30-50 at the time of entry into unemployment. The resulting dataset contains around 4.2 million monthly observations, 435,000 unemployment spells, and 278,000 individuals. Some descriptive statistics are provided in Table 1.

Table 1
Descriptive statistics

	Men	Women
Number of individuals	140,707	138,574
Number of unemployment spells	230,245	204,561
Number of observations	2,070,942	2,201,190
Average duration at spell completion or censoring (months)	8.73	10.38
Average transition rate in first duration month	0.14	0.11
Average transition rate	0.08	0.06
Fraction of individuals with more than one spell	0.37	0.31
Average number of spells for persons with more than one spell	2.72	2.12

[3] Since unemployment insurance is compulsory in Norway, this implies that the analysis is limited to persons who had a job prior to the unemployment spell and who lost this job involuntarily.

3.1 The homogenous cycle model

The *homogenous cycle model* ended up with six support points in the unobserved heterogeneity distribution. The total number of free parameters in this model was 276. Through the process of successive inclusion of unobserved heterogeneity, the log-likelihood increased with 3888 units, from -1021225.91 (without unobserved heterogeneity) to -1017338. We first present the main results of interest, with a particular emphasis of calendar time effects (σ_t) and the resultant labour market tightness indicator. We thereafter discuss how the estimated model can be used for simulation purposes in order to throw light on how employment transition rates - and their evolvement over time and spell durations - are affected by the four main sources of variation, i.e. i) calendar time, ii) spell duration, iii) observed heterogeneity, and iv) unobserved heterogeneity.

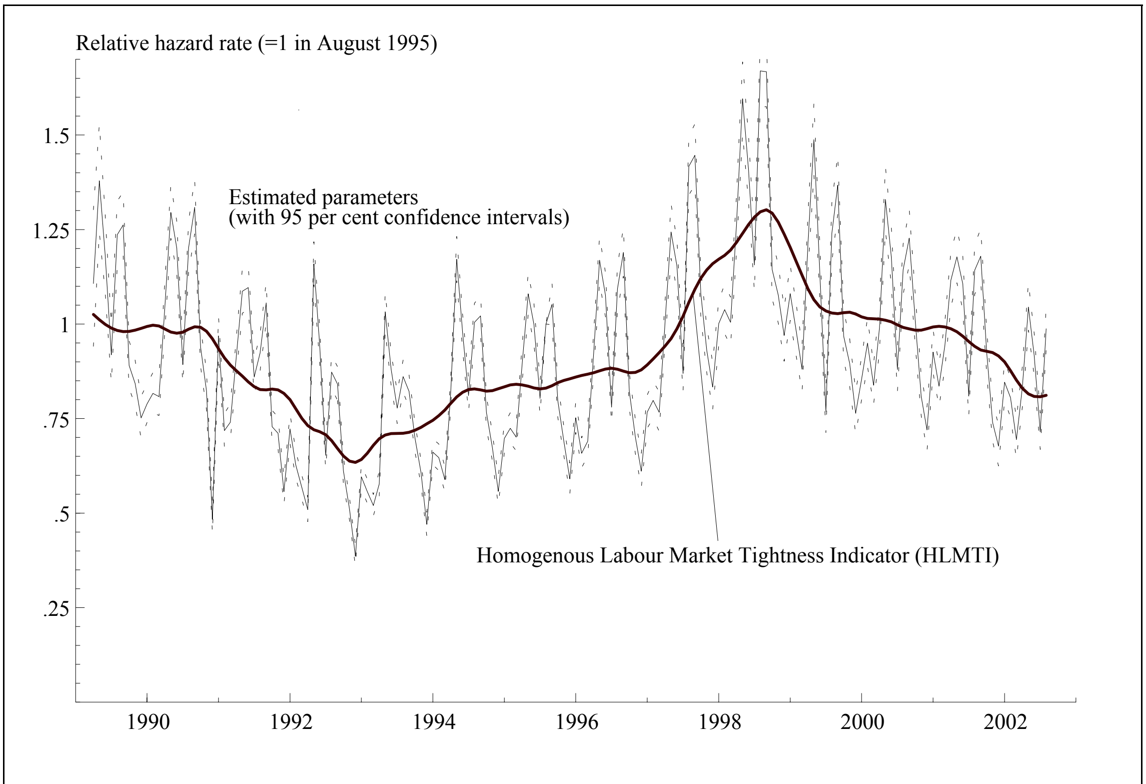

Figure 1 Estimated calendar time parameters and the resulting Homogenous Labour Market Tightness Indicator (HLMTI), 1989.4-2002.8.

Figure 1 presents the estimated calendar time parameters, $\exp(\hat{\sigma}_t)$, divided by the number of trading days in each month (together with 95 per cent point-wise confidence intervals). The confidence intervals are extremely tight; hence from now on we focus on point estimates. There are fairly large high-frequency movements in calendar time effects, and some form of smoothing seems required in order to extract a business cycle pattern. We use the trend-cycle component depicted in Figure 1 as a *Homogenous Labour Market Tightness Indicator (HLMTI)*. This indicator is obtained by applying a standard X12-ARIMA filter (Bureau of the Census, 1999) to the trading-day-adjusted series of $\exp(\hat{\sigma}_t)$. The indicator tells the following story about the Norwegian business cycle pattern during the past 14 years: Labour market prospects deteriorated sharply from 1989 to 1993. In the course of this period, the employment hazard rate declined with approximately 40 per cent, ceteris paribus. The trough occurred in December 1992. From the spring of 1993, employment prospects improved steadily until the autumn of 1998, with a particularly strong improvement in 1997 and in the spring of 1998. During the whole period, the employment hazard rate doubled, ceteris paribus. The peak occurred in September 1998. The 1997-98 business cycle boom was replaced by a sharp decline during the autumn of 1998 and the spring of 1999. From the summer of 1999, there has been a more stable development, but employment prospects have continued to deteriorate. There are some indications that the business cycle downturn gained force once again during the spring of 2002.

In Figure 2, we have plotted our own tightness indicator (HLMTI) from Figure 1, together with two other popular business cycle indicators; the rate of employment (1-the registered rate of unemployment[4]), and the estimated deviation of GDP from its

[4] We use the registered rate of open unemployment, corrected for a break in January 1999.

trend (calculated by Statistics Norway[5]). In order to make the time series properties of the three variables directly comparable, all the series are standardised (i.e. we have subtracted the mean and divided by the standard deviation). The first thing to note is that the hazard-based tightness indicator (HLMTI) matches the GDP-based indicator perfectly with respect to the *timing* of the two business cycle turning points during the 1990's. The main difference seems to be that the GDP-based indicator is smoother, and hence do not capture smaller changes in the business cycle pattern. The second thing to note is that the rate of unemployment apparently behaved pro-cyclically around the first of these turning points (Spring, 1993), and non-cyclically around the second (Autumn, 1998). In both cases, the turning points in the rate of unemployment occurred around six month after the turning points in the underlying transition probability. Hence, to the extent that policy makers founded their assessment of labour market tightness on the development of the unemployment rate, they could have been misguided in both these periods. Moreover, if we look at the recent developments, it seems that the rate of unemployment underestimates the speed by which the Norwegian labour market has slacked during 2001 and 2002.

[5] See Johansen and Eika (2000) for a description of the methodology used to identify the deviation of GDP from its trend.

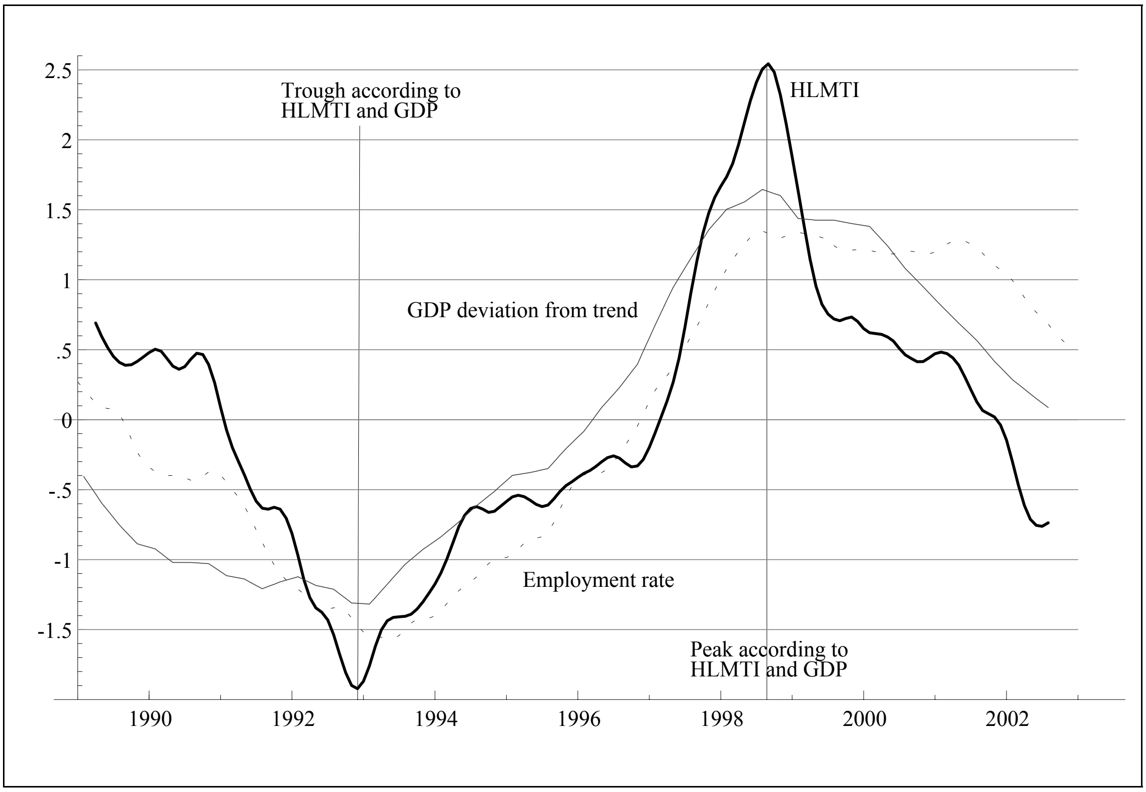

Figure 2. Standardised employment rates and the Homogenous Labour Market Tightness Indicator (HLMTI) (deviations from mean divided by standard deviation)

Policy makers do of course not only base their assessment of labour market developments on the rate of unemployment. But since the unemployment rate is one of the few business cycle statistics that are available almost continuously without delay, it does seem to play a relatively important role in practice. Even though GDP-based measures are typically considered to give a more accurate description of business cycle developments *ex post*, they are of limited value for the assessment of current developments, since GDP numbers are calculated with long delays, and often subject to major revisions. Our own hazard-based tightness indicator seems to combine the virtues of *speed* and *accuracy*; it can be calculated just as quickly as the unemployment rate, and it tracks the performance of the GDP-based measure.

The reason why the rate of unemployment may behave in a pro-cyclical fashion around the business cycle turning points is that it takes some time to level out the flows into and out of unemployment. Exactly how long time it takes obviously depends (among other things) on the composition of the unemployment pool at the different stages of the cycle. At a business cycle trough, there is a relatively large fraction of long-term unemployed. Long-term unemployed tends to have lower employment transition rates than short-term unemployed. In our data, the average transition rate for a person with 12 months of unemployment is 56 per cent lower than for a person with only one month of unemployment. This is to some extent a pure *selection* phenomenon (people become long-term unemployed precisely because they have low individual transition rates). In addition, there may be *structural duration dependence* implying that the length of the spell has a direct causal impact on the employment prospects. The latter of these mechanisms is illustrated in Figure 3, where we have plotted the estimated individual spell duration effects (the degree of structural duration dependence). On average, there seems to be strong negative duration dependence in employment hazards, except during the months just prior to potential benefit termination around the 18th duration month[6]. But, as indicated in the lower panel of Figure 4, the degree of duration dependence vary among different demographic groups. Negative duration dependence is stronger the higher is the age and the better is the education. This probably reflects that longer spells of unemployment are more stigmatising (or demoralising) the older and more educated are the job seekers.

[6] During most of the estimation period, Norwegian labour market authorities practiced a sort of 'soft constraint' on maximum benefit duration. The maximum duration was formally 80 weeks (approximately 18 months), but with ample scope for exemptions and renewals (see Røed and Zhang, 2003).

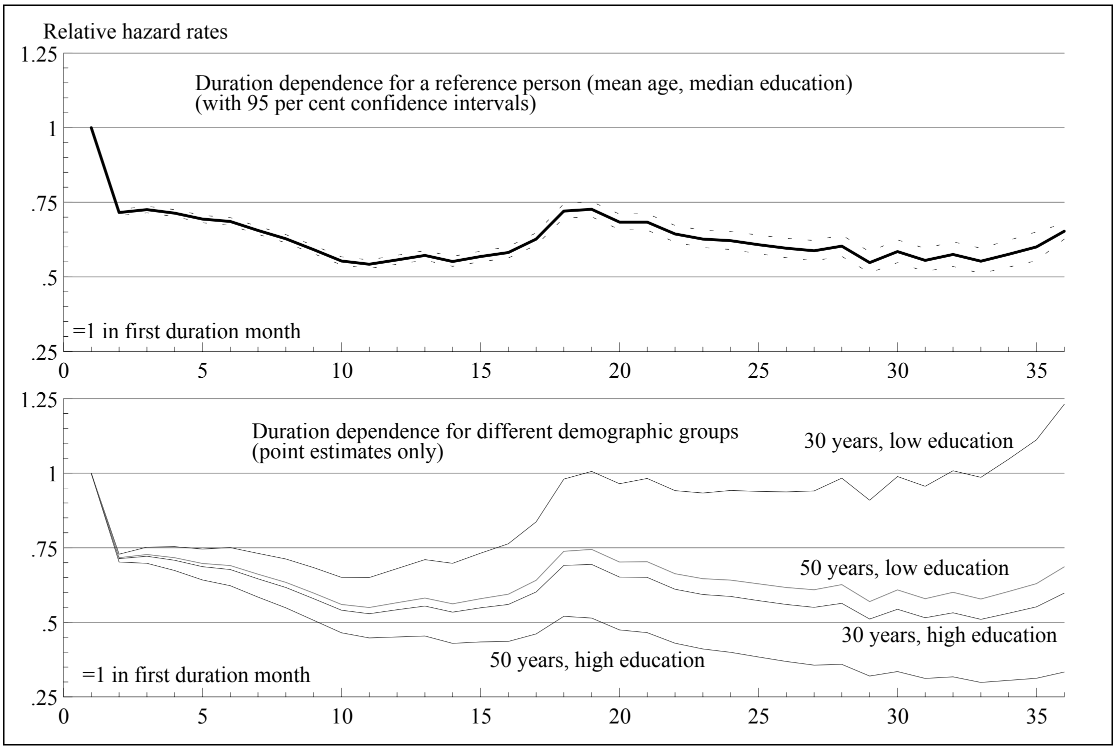

Figure 3. Estimated duration dependence during the first three years of unemployment
Note: The spell duration pattern is estimated with the aid of one indicator variable for each possible duration plus linear interaction terms between duration and age and duration and educational attainment.

We now take a closer look at the *selection processes* among the unemployed – and how these processes interact with calendar time and spell duration, by using the estimated model to simulate the progression of our unemployment spells, *given their actual starting dates*. Figure 4 depicts the actual monthly transition rates from unemployment to employment, together with the predicted transition rates based on the model simulation. The main reason why we do not get a perfect match is that there are a number of time-varying covariates in the actual data, and also a substantial number of spells that were censored due to benefit exhaustion, disability etc., which we have not been able to take properly into account in the simulation exercise. The simulated transition probabilities can nevertheless be used as a tool for decomposition with respect to the different sources of variation.

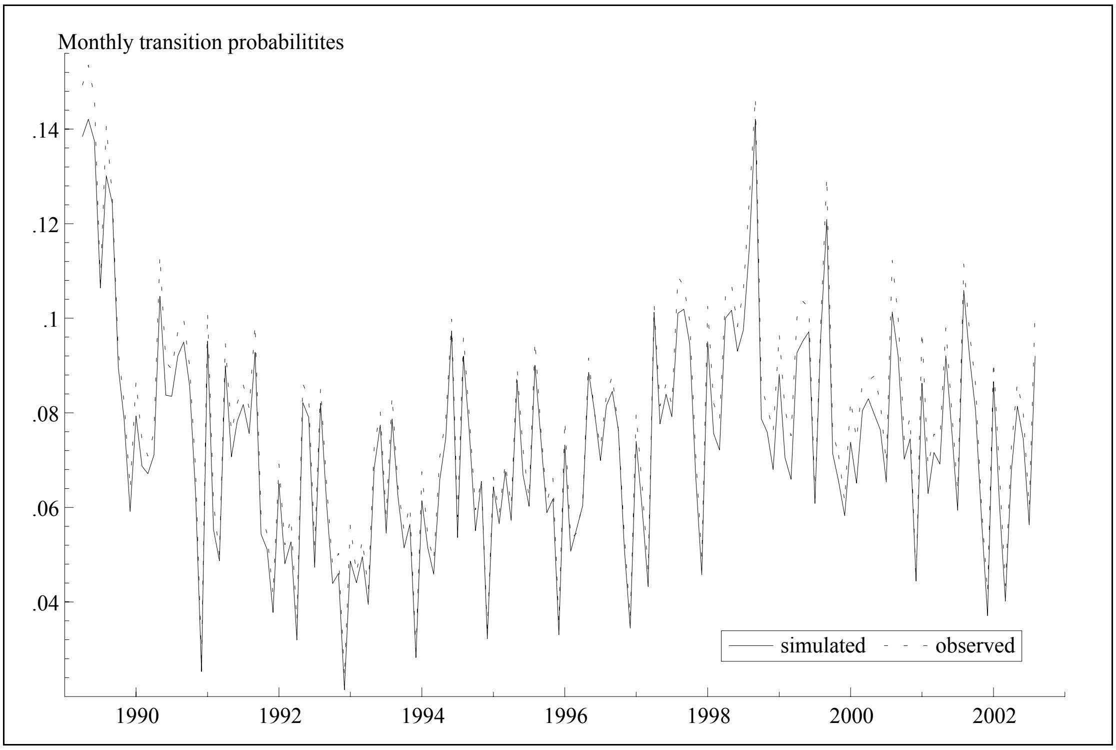

Figure 4. Observed and simulated monthly transition rates 1989.4-2002.8

The simulated data contains a predicted employment transition probability for each observation month, calculated on the basis of the four sources of variation in these probabilities; calendar time, spell duration, observed heterogeneity and unobserved heterogeneity. Hence, it is possible to decompose the overall variation in the transition rates into the various sources. Given the large calendar time variation in employment transition rates that we have revealed, one should perhaps expect to find that calendar time effects account for a very large fraction of the variation in the employment transition rates. But this is not the case. The total variance of predicted employment transition rates ($\mathrm{var}\,\hat{\varphi}$) can be decomposed into its within-month and across-months components in the following way: $\mathrm{var}\,\hat{\varphi} = E\left[\mathrm{var}\,\hat{\varphi} \mid t\right] + \mathrm{var}\left[E\left[\hat{\varphi} \mid t\right]\right]$. According to this decomposition, the across-month component accounts for only 12 per cent of the transition rate variance in our simulated data. Hence, there is a substantial element of within-month heterogeneity in individuals' employment transition rates. This heterogeneity is, however, relatively stable over time. Figure 5 illustrates this point. The upper left-hand panel repeats the predictions from Figure 4, while the three other panels depicts predicted transition rates after the removal of the estimated effects of calendar time, spell duration and observed heterogeneity (in a cumulative fashion). The removal of calendar time effects does away with most of the calendar time variation in the average predicted hazard rates. The remaining variation is due to different types of selection mechanisms that are correlated with calendar time. In particular, there is a visible tendency for transition rates to be high during the fist 12-18 months of the estimation period, reflecting that our flow-sampling scheme makes short-term unemployed strongly over-represented in this period. Apart from this, there seems to be relatively small compositional changes in the unemployment pool over time, at least compared to the calendar time effects themselves. But the selection effects are not completely absent. This can be seen by taking a closer look at the lower right-hand-side panel in Figure 5. In Figure 6, we look at the pure selection effects due to unobserved heterogeneity in transition rates over time in the period from 1992-2002 with a much finer scale on the vertical axis. The selection effects in the calendar time pattern of transition rates then become clearly visible. Average unobserved 'employability' among the unemployed deteriorated steadily until the autumn of 1998, but these composition effects did not reduce the average monthly employment transition rate by more than 0.5-0.6 percentage points. It is also evident that these compositional changes lag the business cycle with several years, reflecting that e.g. an economic upturn initially may promote, rather than counteract, selection mechanisms, as the most employable job seekers are the first to take advantage of the improved job prospects.

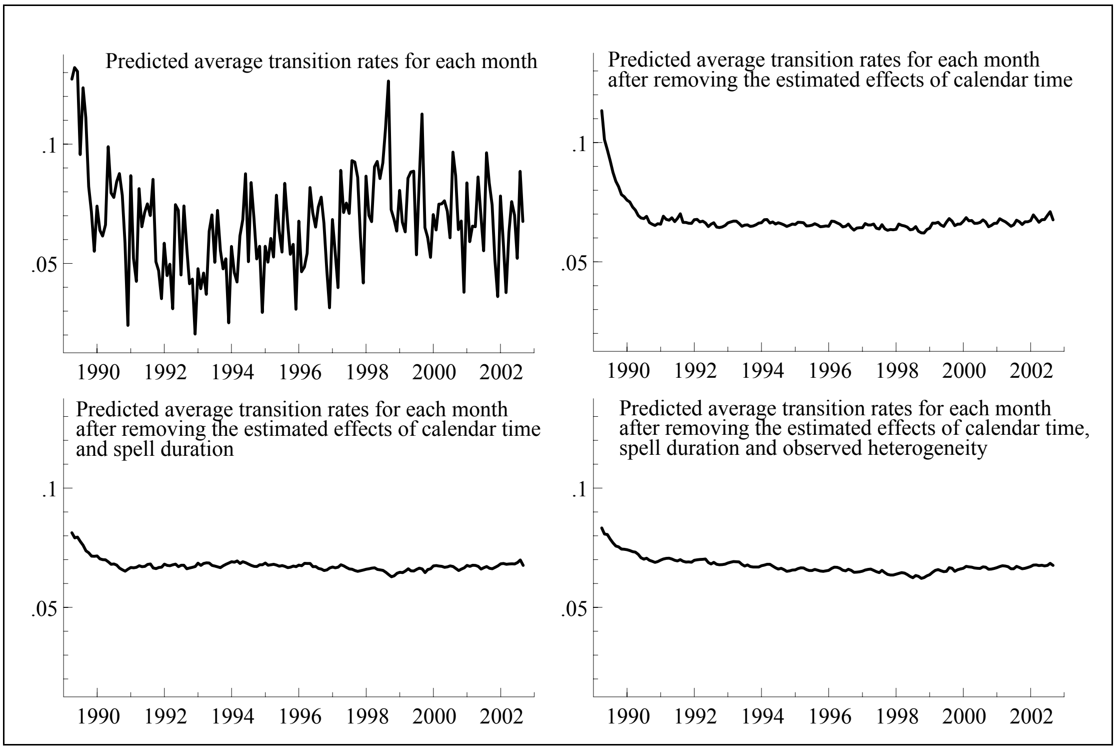

Figure 5. Average predicted hazard rates 1989.4-2002.8, based on simulated data.
Note: The various sources of variation are removed by replacing the actual values of the relevant variables with a constant that keeps the overall average unchanged.

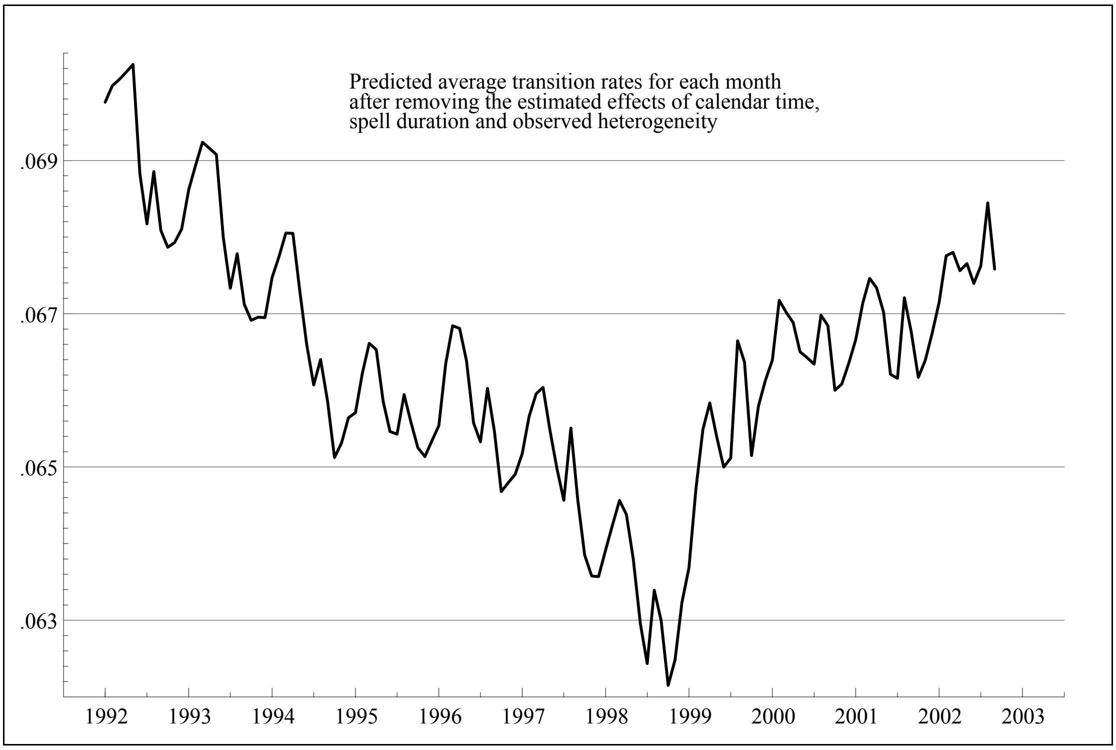

Figure 6. A closer look at average predicted hazard rates 1992.1-2002.8, based on the role of unobserved heterogeneity only.
Note: This figure is the same as the figure in the lower right-hand-side panel in Figure 5, only with a different time period and axis.

The driving force behind the compositional changes in employability over time is that the spell-duration composition also changes over time. There are strong selection effects with respect to spell duration. Figure 7 disentangles the different mechanisms that are responsible for creating a negative correlation between the employment transition rate and spell duration. The upper left hand panel depicts how the average predicted transition rates developed over spell duration during the estimation period. In the upper right-hand panel, we see that the removal of calendar time effects hardly affect this relationship at all. When the estimated spell duration effects are removed, the pattern of course changes dramatically. But a strong negative correlation remains, and it can clearly be seen from the lower right-hand panel, that most of this arises from unobserved heterogeneity[7]. Unobserved heterogeneity *alone* is responsible for reducing the observed employment transition rate with approximately 20 per cent from the first to the sixth duration month.

[7] By comparing the two lower panels, it can be seen that observed heterogeneity to some extent induces a positive correlation between transition rates and spell duration, particularly from the first to the second duration month. One reason for this is that there are some important time-varying covariates, reflecting participation in labour market programs and access to some part-time work, that as a matter of definition are equal to zero in the first duration month (since we have conditioned on full time open unemployment in order to start a new spell).

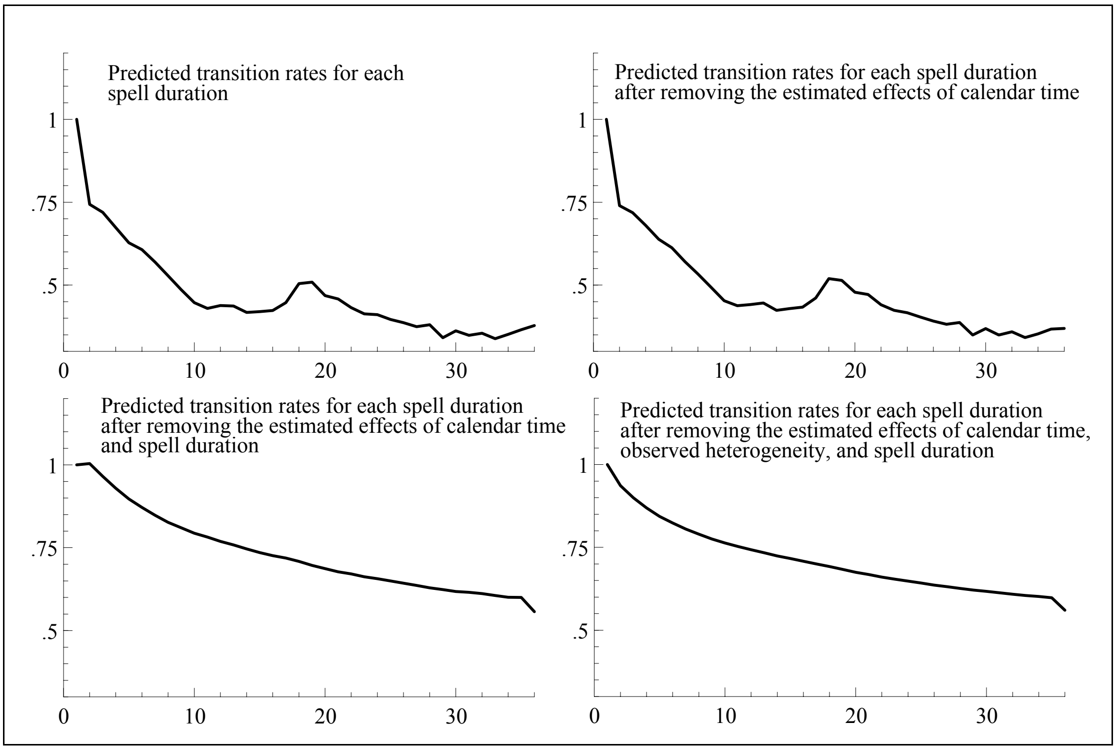

Figure 7 Average predicted hazard rates for each spell duration from 1 to 36 months, based on simulated data.
Note: The last month transition rate covers transition at 36 month duration *or higher*. The various sources of variation are removed by replacing the actual values of the relevant variables with a constant that scales the first month relative transition rate to unity.

The apparently weak tendency for selection mechanisms to affect the calendar time patterns of employment transition rates suggests that the crude transition rate from unemployment to employment, may serve as the basis for calculation of labour market tightness indicators. And indeed, as it turns out, our homogenous tightness indicator (HLMTI) does not differ very much from the crude transition rate in the data (smoothed by X12-ARIMA). Figure 8 illustrates this point. There are substantial differences between HLMTI and the crude outflow rate in the beginning of the estimation period, but again, this is only an artefact of our *flow based* sampling with a strong over-representation of short spells in the beginning of the period. As the distribution of spell lengths becomes more representative, the two series trace each other quite closely, although there is a weak tendency for the crude outflow rate to underestimate the effects of business cycle changes. This suggests that the *outflow rate from unem-*

ployment (properly adjusted for e.g. seasonality and trading days) may in practice serve as a quite reliable labour market tightness indicator. However, it is difficult to assess the generality of this result. Compared to many other European economies, Norway is characterised by a relatively low unemployment rate, with few really long-term unemployed (particularly in the core group of prime aged previously employed unemployed that are used in the present analysis). It is likely that the selection issue is more important in labour markets with a very large fraction of long-term unemployed, and in which some of these long-term unemployed in reality have a zero employment hazard.

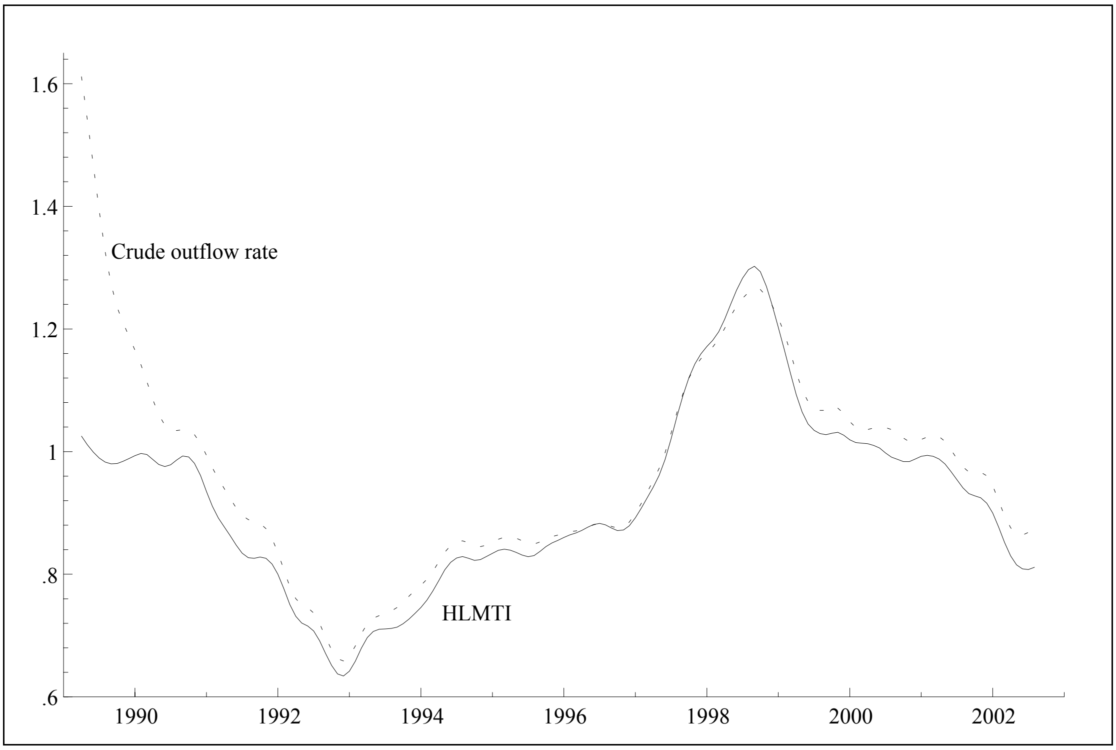

Figure 8. The homogenous labour market tightness indicator (HLMTI) and the trend cycle component of the crude outflow rate (X12ARIMA).

3.2 Idiosyncratic labour market tightness indicators

So far, we have assumed that the business cycle pattern is the same for all unemployed persons. This need not be the case. Idiosyncratic business cycle patterns may arise along several dimensions, e.g. related to personal characteristics, occupation, spell duration or region. In this section, we give a few examples of how the hazard rate approach can be exploited to shed light on business cycle developments for different types of job seekers separately. We first divide the set of observations into 12 groups according to gender, educational attainment (two groups) and spell duration (three groups), and estimate a separate business cycle indicator for each of them. Even with this relatively modest grouping exercise, we obtain as much as 1,910 calendar time parameters and 2,020 free parameters in total[8]. The main results from this exercise is presented in Figure 9 in the form of standardised business cycle indicators for each of the 12 groups. The bottom line seems to be that business cycle developments have produced similar time patterns in the hazard rates for the various groups. In particular, the *timing* of the two turning points is almost the same for all groups. The main difference seems to be that, compared to the cyclical slump in the early 1990's, the most recent economic downturn has had a relatively much stronger negative effect on persons with high education than on persons with low education (particularly at short durations).

[8] We restricted the heterogeneity distribution to contain two points of support in this case. The log-likelihood obtained was -1013636.69.

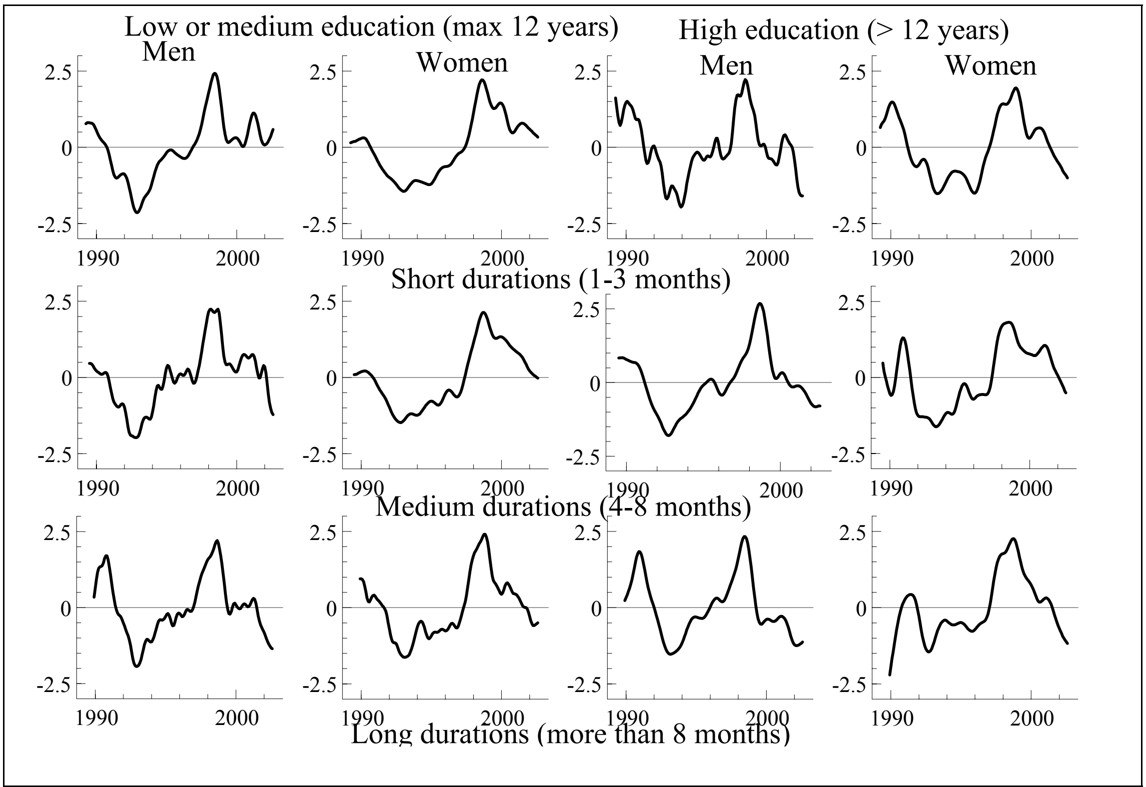

Figure 9. Standardised idiosyncratic labour market tightness indicators for 12 observation groups, divided according to gender, education and spell duration.
Note: The series are standardised by subtracting their means and dividing by their standard deviations.

For labour market authorities, it may be more important to obtain knowledge about the business cycle developments for different occupations or industries. For that purpose, we divided the population into six groups, based on their previous work experience. This model ended up 960 calendar time parameters and 1078 parameters in total[9]. The main results are illustrated in Figure 10. Again, we find that the timing of turning points is quite similar across the various groups, but that the relative strength of the two 'recessions' (the one in the early 1990's and the one in the early 2000's) differs.

[9]The log-likelihood obtained its maximum of -1013785.30 with seven support points in this case.

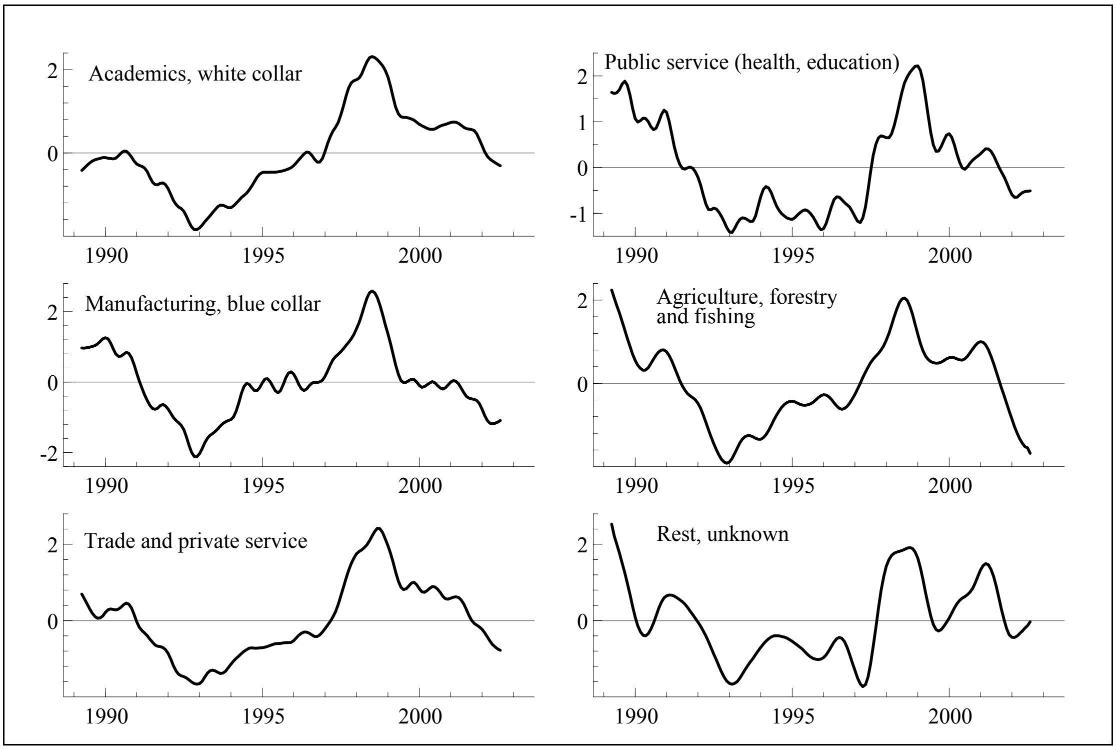

Figure 10. Standardised idiosyncratic labour market tightness indicators for 6 observation groups, divided according to previous occupation.
Note: The series are standardised by subtracting their means and dividing by their standard deviations.

In addition to the above examples, it may be of particular interest to compute *regional* business cycle indicators. Such indicators can be used as an alternative to regional unemployment rates, both in order to assess local labour market developments i more detail, and to account for labour market tightness in e.g. models of local wage formation and regional migration. The local rate of unemployment is well known to play a prominent role in the latter types of models, and a relationship between the wage level and the local rate of unemployment has even been elevated into an empirical law (Blanchflower and Oswald, 1994). But since the rate of unemployment in these models only serve as a proxy for the underlying transition rates, one should perhaps expect that both wage formation and migration behaviour could be more accurately predicted, and also better understood, if the rate of unemployment was replaced (or accompanied) by a labour market tightness indicator of the kind that we have derived above. We do not pursue this idea in the present paper. But the issue is investigated in Carlsen et al (2003). In that paper, Norwegian register data is used to estimate a *yearly* labour market tightness indicator (based on employment transition rates during each month in each year) for 90 Norwegian regions during the 1990's. These indicators are then included, together with alternative local labour market tightness measures (vacancy -and unemployment rates), in econometric models that explain local wage formation as well as regional migration during this time period. The results are striking: The idiosyncratic hazard-based tightness indicator outperforms the rate of other indicators, and renders them superfluous in models of both wage formation and migration.

4 Concluding Remarks

In this paper we have shown that the rate of unemployment is a poor indicator for business cycle developments, particularly around business cycle turning points. We have derived an alternative unemployment-based tightness indicator, which is essentially a heterogeneity-adjusted outflow rate from the unemployment pool. We have demonstrated that this indicator displays its turning points in accordance with a GDP-based business cycle indicator. The heterogeneity adjustment turned out to have had a relatively modest impact on the time series properties of the outflow rate in Norway during the 1989-2002 period. Hence, it seems that the crude outflow rate from the pool of unemployment may serve as a remarkably accurate indicator of labour market tightness, as long as the outflow rate is calculated on the basis of the 'core' group in the labour market (i.e. prime aged benefit claimants).

Our results indicate that labour market authorities and statistical agencies should put more emphasis into the production of reliable flow-statistics, and that policy makers should keep a close eye on these statistics' development. A rise (fall) in the properly smoothed exit rate from unemployment 'now' seems to be a strong signal indicating that the rate of unemployment will start to fall (rise) within a period of 6-12 months.

Appendix: Maximisation of the Likelihood and the Concept of Implicit Dummy Variables

The models presented in this paper cannot be estimated with 'standard' software. The reason is that the amount of data and the number of parameters required to estimate non-parametric calendar time effects precisely (without unjustified restrictions on e.g. unobserved heterogeneity or spell duration effects), produces an optimisation problem that exceeds normal computational capacity by several orders of magnitude. We have therefore constructed our own software package. In this Appendix, we briefly present the main content of this package.

The non-linear maximisation method itself is standard; we currently use an initial step of BFGS (the LBFGS package by Liu & Nocedal) and switches to a Newton method (a modification of Xie and Schlick's TNPACK with analytic Hessian) after a while. For the Newton method we use Fisher's matrix as an approximation to the Hessian. Applied on the kind of functions and data used in this paper, almost all process time is spent on computation of the likelihood and its gradient. Comparatively little time is spent in the internal arithmetic of the BFGS and Newton method.

The essence of the non-parametric approach is that the number of explanatory variables is extremely large, but that most of them are binary indicator variables (dummy variables). For example, in order to estimate the effects of calendar time *without parametric restrictions*, we must attribute one parameter to each calendar time unit that occurs in the dataset. In our case, this amounts, for example, to estimate approximately 160 calendar month parameters (p1,...,p160), one for each calendar month dummy variable (month1,...,month160). If these dummy variables are treated just like ordinary explanatory variables, the computational task quickly becomes unmanageable. Every time the likelihood is computed, we would evaluate the sum (*p1*month1+...+ p160*month160*). This amounts to 160 multiplications and 159 additions, together with 320 memory lookups. But since all the dummies are zero except for one, the sum is just, say, *p78* (if 78 is the value of the original variable *month*). That is, everything is zero except for one term in the sum. In contrast to human arithmetic, a computer uses equally long time to multiply and add zeros as it uses to multiply and add anything else. It is therefore more efficient to tell the computer that most terms are zero and let it pick the right parameter directly. This is exactly what an implicit dummy does. We retain the original variable *month*, we do not create any dummy variables *month1* etc, but rather we specify to the computer program that the single variable *month* is an *implicit dummy* with range 1...160. The program then creates a set of parameters *p1... p160* as above and uses the value of *month* to directly lookup the right parameter. In this case, the result is that 329 arithmetic operations and 320 memory lookups is replaced with 2 memory lookups (one for *month* and one for the relevant *parameter*). The mathematical result is exactly the same as if we had created ordinary dummy variables. More importantly, we also utilize the implicit dummies when we compute the gradient. The speedup for the gradient computation is comparable to the speedup for the likelihood function as can be seen from the following elementary calculation: Let $f(x)$ be the contribution to the likelihood function from a single observation, where $x = \sum p_i v_i$ for p_i, v_i parameters and variables. We then have $\frac{df}{dp_i} = \frac{df}{dx}\frac{dx}{dp_i} = v_i \frac{df}{dx}$, where v_i is 0 for all i except one, so that in addition to the speedup in computing x, we also avoid computing and storing a lot of derivatives in the gradient. The potential for speedup in computing the Hessian is even bigger because entire blocks of the matrix are zero. Note, however, that except in special cases, we have not succeeded in getting the Hessian computation numerically stable. Instead we use the Fisher matrix, which is easy to compute from the gradient (it also has the additional advantage of being definite)

The overall speedup of the maximisation is perhaps not as large as the above figures indicate; the reason is that there is more to likelihood maximisation than the arithmetic described above. Whereas the use of implicit dummies effectively may reduce 160 variables to 1, it does not reduce the number of parameters. Also, there is some overhead in using implicit dummy variables; hence it does not pay unless the range of the dummy is more than 4-6.

In order to compare the speed of our own program with that of existing software, we estimated the simplest version of model described in Section 3 (the homogenous cycle model without unobserved heterogeneity) with STATA 7.0 (complemetery log-log command), as well as with our own program. For this purpose, we used a Dell Precision 620, with 1 CPU (800 Mhz) and 2 GB memory. STATA required around 10 hours and 10 minutes estimation time on this machine. Our own program used 4 minutes and 16 seconds. The estimation becomes much more computationally demanding when we add unobserved heterogeneity into the model. As a rule of thumb, the computational cost increases by a factor equal to the average spell-length. Adding additional support points in the heterogeneity distribution increases the complexity linearly in the number of points. In order to carry out the whole estimation procedure described in Section 2 (continue with additional support points until it is no longer possible to improve the likelihood), computation time typically has to be multiplied by a factor of several hundred. Another complicating factor is that when increasing the number of support points, our Fisher-approximation to the Hessian gradually deteriorates. This does not affect the result, but it does affect the convergence speed; sometimes we need more than 50 iterations instead of 7-8.

Even with the most efficient optimisation procedures, it is therefore a formidable computational task to maximise the type of likelihood functions discussed in this paper. In practice, the non-concavity of the likelihood function may imply that the likelihood function must be maximised repeatedly in order to verify that the obtained maximum is really a global one. It may therefore sometimes be desirable to use several computers simultaneously in the optimisation process. Fortunately, the optimisation problem to be solved turns out to be what is known in High Performance Computing circles as being *embarrassingly parallel.* The log-likelihood is a sum over all the observations. The program splits the data up in, say, 16 equally sized parts. Each part is sent to a separate computer (or CPU) for computation and the results are collected and summed afterwards. The speedup is very close to the number of CPUs we run on. We have successfully run the program in parallel on 40 CPUs on a HP Super-Dome computer at the High Performance Computing facility at the University of Oslo, as well as on a cluster of 15 dual-cpu AMD PCs running Linux. There is, however, a limit to how many CPUs we may efficiently use, one reason is that it becomes increasingly hard to distribute the data evenly between the CPUs.

References

Abbring, J. H., Van den Berg, G. J. and Van Ours, J. C. (2002) The Anatomy of Unemployment Dynamics. *European Economic Review*, Vol. 46, 1785-1824.

An, M. Y. (1998) Logconcavity versus Logconvexity: A Complete Characterization. *Journal of Economic Theory*, Vol. 80, 350-369.

Blanchflower D. and Oswald, A. (1994) The Wage Curve. MIT Press, Cambridge MA.

Bover, O., Arellano, M. and Bentolila, S. (2002) Unemployment Duration, Benefit

Duration and the Business Cycle. *Economic Journal*, Vol. 112, No. 479, 223-265.

Brinch, C. (2000) Identification of Structural Duration Dependence and Unobserved Heterogeneity with Time-varying Covariates. Memorandum No. 20/2000, Department of Economics, University of Oslo (www.oekonomi.uio.no/memo/memopdf/memo2000.pdf)

Burdett, K. and Ondrich, J. I. (1985) How Changes in Labor Demand Affect Unemployed Workers. *Journal of Labor Economics*, Vol. 3, No. 1, 1-10.

Bureau of the Census (1999). *X-12-ARIMA Reference Manual*. Bureau of the Census, Time Series Staff, Statistical Research Division, Washington DC.

Carlsen, F., Johansen, K. and Røed, K. (2003) Wage Formation, Migration, and Local Labour Market Tightness. Working Paper Series No. 3/2002, Department of Economics, Norwegian University of Science and Technology (www.svt.ntnu.no/iso/wp/wp.htm).

Heckman, J. and Singer, B. (1984) A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data. *Econometrica*, Vol. 52, 271-320.

Imbens, G. and Lynch, L. M. (1993) Re-Employment Probabilities over the Business Cycle. NBER Working Paper No. 4585.

Johansen, P. R. and Eika, T. (2000) Driving Forces Behind Business Cycles in the 1990's. *Economic Surveys* No. 3/2000, Statistics Norway (http://www.ssb.no/english/subjects/08/05/10/es/200003/johansen.pdf).

Layard, R., Nickell. S and Jackman, R. (1991) Unemployment. Macroeconomic Performance and the Labour Market. Oxford University Press.

Lindsay, B. G. (1983) The Geometry of Mixture Likelihoods: A General Theory. *The*

Annals of Statistics, Vol. 11, 86-94.

Røed, K. and Raaum, O. (2003) Administrative Registers – Unexplored Reservoirs of Scientific Knowledge? *Economic Journal*, forthcoming.

Røed, K. and Zhang, T. (2003) Does Unemployment Compensation Affect Unemployment Duration? *Economic Journal*, Vol. 113, 190-206.

Van den Berg, G. J. (2001) Duration Models: Specification, Identification, and Multiple Durations. *Handbook of Econometrics*, Vol. 5, Chapter 55, Elsevier.